

Global Alignment of Molecular Sequences via Ancestral State Reconstruction

Alexandr Andoni¹ Constantinos Daskalakis² Avinatan Hassidim³ Sebastien Roch⁴

¹Department of Computer Science, Princeton University, Princeton, NJ 08544

²CSAIL, MIT, Cambridge, MA 02139

³MIT, Cambridge, MA 02139

⁴Department of Mathematics and Bioinformatics Programs, UCLA, Los Angeles, CA 90024

andoni@mit.edu costis@csail.mit.edu avinatan@mit.edu roch@math.ucla.edu

Abstract:

Molecular phylogenetic techniques do not generally account for such common evolutionary events as site insertions and deletions (known as indels). Instead tree building algorithms and ancestral state inference procedures typically rely on substitution-only models of sequence evolution. In practice these methods are extended beyond this simplified setting with the use of heuristics that produce global alignments of the input sequences—an important problem which has no rigorous model-based solution. In this paper we open a new direction on this topic by considering a version of the multiple sequence alignment in the context of stochastic indel models. More precisely, we introduce the following *trace reconstruction problem on a tree* (TRPT): a binary sequence is broadcast through a tree channel where we allow substitutions, deletions, and insertions; we seek to reconstruct the original sequence from the sequences received at the leaves of the tree. We give a recursive procedure for this problem with strong reconstruction guarantees at low mutation rates, providing also an alignment of the sequences at the leaves of the tree. The TRPT problem without indels has been studied in previous work (Mossel 2004, Daskalakis et al. 2006) as a bootstrapping step towards obtaining information-theoretically optimal phylogenetic reconstruction methods. The present work sets up a framework for extending these works to evolutionary models with indels.

In the TRPT problem we begin with a random sequence x_1, \dots, x_k at the root of a d -ary tree. If vertex v has the sequence y_1, \dots, y_{k_v} , then each one of its d children will have a sequence which is generated from y_1, \dots, y_{k_v} by flipping three biased coins for each bit. The first coin has probability p_s for Heads, and determines whether this bit will be substituted or not. The second coin has probability p_d , and determines whether this bit will be deleted, and the third coin has probability p_i and determines whether a new random bit will be inserted. The input to the procedure is the sequences of the n leaves of the tree, as well as the tree structure (but not the sequences of the inner vertices) and the goal is to reconstruct an approximation to the sequence of the root (the DNA of the ancestral father). For every $\chi > 0$ we present an algorithm which outputs with probability $1 - \chi$ an approximation of x_1, \dots, x_k if $p_i + p_d < O(1/k^{2/3} \log n)$ and $(1 - 2p_s)^2 > Cd^{-1} \log d$ for some constant $C > 0$, and every large enough d .

To our knowledge, this is the first rigorous trace reconstruction result on a tree in the presence of indels.

Keywords: Evolutionary trees; ancestral reconstruction; sequence alignment; trace reconstruction

1 Introduction

Trace reconstruction on a star. In the “trace reconstruction problem” (TRP) [1–6], a random binary string X of length k generates an i.i.d. collection of traces Y_1, \dots, Y_n that are identical to X except for random *mutations* which consist in *indels*, i.e., the deletion of an old site or the insertion of a new site between existing sites, and *substitutions*, i.e., the flipping of the state at an existing site. (In keeping with biological terminology, we refer to the components or positions of a string as *sites*.) The goal is to recon-

struct efficiently the original string with high probability from as few random traces as possible.

An important motivation for this problem is the reconstruction of ancestral DNA sequences in computational biology [3, 4]. One can think of X as a gene in an (extinct) ancestor species 0. Through speciation, the ancestor 0 gives rise to a large number of descendants $1, \dots, n$ and gene X evolves independently through the action of mutations into sequences Y_1, \dots, Y_n respectively. Inferring the sequence X of an ancient gene from extant descendant copies Y_1, \dots, Y_n is a

standard problem in evolutionary biology [7]. The inference of X typically requires the solution of an auxiliary problem, the *multiple sequence alignment problem* (which is an important problem in its own right in computational biology): site t_i of sequence Y_i and site t_j of sequence Y_j are said to be *homologous* (in this simplified TRP setting) if they descend from a common site t of X *only through substitutions*; in the multiple sequence alignment problem, we seek roughly to uncover the homology relation between Y_1, \dots, Y_n . Once homologous sites have been identified, the original sequence X can be estimated, for instance, by site-wise majority.

The TRP as defined above is an *idealized* version of the ancestral sequence reconstruction problem in one important aspect. It ignores the actual phylogenetic relationship between species $1, \dots, n$. A *phylogeny* is a (typically, binary) tree relating a group of species. The leaves of the tree correspond to extant species. Internal nodes can be thought of as extinct ancestors. In particular the root of the tree represents the most recent common ancestor of all species in the tree. Following paths from the root to the leaves, each bifurcation indicates a speciation event whereby two new species are created from a parent. An excellent introduction to phylogenetics is [8].

A standard assumption in computational phylogenetics is that genetic information evolves from the root to the leaves according to a Markov model on the tree. Hence, the stochastic model used in trace reconstruction can be seen as a special case where the phylogeny is *star-shaped*. (The substitution model used in trace reconstruction is known in biology as the Cavender-Farris-Neyman (CFN) [9–11] model.) It may seem that a star is a good first approximation for the evolution of DNA sequences. However extensive work on the so-called “reconstruction problem” in theoretical computer science and statistical physics has highlighted the importance of taking into account the full tree model in analyzing the reconstruction of ancestral sequences.

The “reconstruction problem.” In the “reconstruction problem” (RP), we have a single site which evolves through substitutions only from the root to the leaves of a tree. In the most basic setup which we will consider here, the tree is d -ary and each edge is an independent symmetric indel-free channel where the probability of a substitution is a constant $p_s > 0$. The goal is to reconstruct the state at the root given the vector of states at the leaves. More generally, one can consider a sequence of length k at the root where each site evolves independently according to the Markov process above. Denote by n the number of leaves in

the tree. The RP has attracted much attention in the theoretical computer science literature due to its deep connections to computational phylogenetics [12–15] and statistical physics [16–27]. See e.g. [28, 29] for background.

Unlike the star case, the RP on a tree exhibits an interesting thresholding effect: on the one hand, information is lost at an exponential rate along each path from the root; on the other hand, the number of paths grows exponentially with the number of levels. When the substitution probability is low, the latter “wins” and vice versa. This “phase transition” has been thoroughly analyzed in the theoretical computer science and mathematical physics literature—although much remains to be understood. More formally, we say that the RP is *solvable* when the correlation between the root and the leaves persists no matter how large the tree is. Note that unlike the TRP we do not require high-probability reconstruction in this case as it is not information-theoretically achievable for d constant—simply consider the information lost on the first level below the root. Moreover the “number of traces” is irrelevant here as it is governed by the depth of the tree and the solvability notion implies nontrivial correlation for any depth. When the RP is unsolvable, the correlation decays to 0 for large trees. The results of [17, 20, 22, 23, 30, 31] show that for the CFN model, if $p_s < p^*$, then the RP is solvable, where $d(1 - 2p^*)^2 = 1$. This is the so-called *Kesten-Stigum* bound [32]. If, on the other hand, $p_s > p^*$, then the RP is *unsolvable*. Moreover in this case, the correlation between the root state and any function of the character states at the leaves decays as $n^{-\Omega(1)}$. The positive result above is obtained by taking a majority vote over the leaf states.

Like the TRP, the RP is only an *idealized* version of the ancestral sequence reconstruction problem: it ignores the presence of indels. In other words, the RP assumes that the multiple sequence alignment problem has been solved perfectly. This is in fact a long-standing assumption in evolutionary biology where one typically preprocesses sequence data by running it through a multiple sequence alignment heuristic and then one only has to model the substitution process. This simplification has come under attack in the biology literature, where it has been argued that alignment procedures often create systematic biases that affect analysis [33, 34]. Much empirical work has been devoted to the proper joint estimation of alignments and phylogenies [33, 35–41].

Our results. We make progress in this recent new direction by analyzing the RP in the presence of indels—which we also refer to as the TRP on a tree

(TRPT). We consider a d -ary tree where each edge is an independent channel with substitution probability p_s , deletion probability p_d , and insertion probability p_i (see Section 1.1 for a precise statement of the model). The root sequence has length k and is assumed to be uniform in $\{0, 1\}^k$. As in the standard RP, we drop the requirement of high-probability reconstruction and seek instead a reconstructed sequence that has correlation with the true root sequence uniformly bounded in the depth.

We give an efficient recursive procedure which solves the TRPT for $p_s > 0$ a small enough constant (strictly below, albeit close, to the Kesten-Stigum bound) and $p_d, p_i = O(k^{-2/3} \log^{-1} n)$. As a by-product of our analysis we also obtain a partial global alignment of the sequences at the leaves. Our method provides a framework for separating the indel process from the substitution process by identifying well-preserved subsequences which then serve as markers for alignment and reconstruction (see Section 1.2 for a high-level description of our techniques). As far as we are aware, our results are the first rigorous results for this problem.

Results on the RP have been used in previous work to advance the state of the art in rigorous phylogenetic tree reconstruction methods [13–15, 42]. A central component in these methods is to solve the RP on a partially reconstructed phylogeny to obtain sequence information that is “close” to the evolutionary past; then this sequence information is used to obtain further structural information about the phylogeny. The whole phylogeny is built by alternating these steps. Our method sets up a framework for extending these techniques beyond substitution-only models.

Related work. Much work has been devoted to the trace reconstruction problem on a star [1–6]. In particular, in [5], it was shown that, when there are only deletions, it is possible to tolerate a small constant deletion rate using $\text{poly}(k)$ traces. For a different range of parameters, Viswanathan and Swaminathan [6] showed that, under constant substitution probability and $O(1/\log k)$ indel probability, $O(\log k)$ traces suffice. Both results assume that the root sequence X is uniformly random.

The multiple sequence alignment problem as a combinatorial optimization problem (finding the best alignment under a pairwise scoring function) is known to be NP-hard [43, 44]. Most heuristics used in practice, such as CLUSTAL [45], T-Coffee [46], MAFFT [47], and MUSCLE [48], use the idea of a guide tree, that is, they first construct a very rough phylogenetic tree from the data (using edit distance as a measure of evolutionary distance), and then

recursively construct local alignments produced by “aligning alignments.” Our work can be thought as an attempt to analyze rigorously this type of procedure. Note that the Steiner version of the multiple sequence alignment problem on a fixed phylogeny, the so-called *tree alignment* problem, is known to admit a polynomial-time approximation scheme [49, 50].

Finally, our work is tangentially related to the study of edit distance. Edit distance and pattern matching in random environments has been studied, e.g., by [51–53].

1.1 Definitions

We now define our basic model of sequence evolution.

Definition 1.1 (Model of sequence evolution)

Let $T_H^{(d)}$ be the d -ary tree with H levels and $n = d^H$ leaves. For simplicity, we assume throughout that d is odd. We consider the following model of evolution on $T_H^{(d)}$. The sequence at the root of $T_H^{(d)}$ has length k and is drawn uniformly at random over $\{0, 1\}^k$. Along each edge of the tree, each site (or position) undergoes the following mutations independently of the other sites:

- **Substitution.** The site state is flipped with probability $p_s > 0$.
- **Deletion.** The site is deleted with probability $p_d > 0$.
- **Insertion.** A new site is created to the right of the current site with probability $p_i > 0$. The state of this new site is uniform $\{0, 1\}$.

These operations occur independently of each other. The last two are called indels. We let $p_{id} = p_i + p_d$ and $\theta_s = 1 - 2p_s$. The parameters p_s, p_d, p_i may depend on k and n , where n is the number of leaves.

1.2 Results

Statement of results. Our main result is the following. Denote by $X = x_1, \dots, x_k$ a binary uniform sequence of length k . Run the evolutionary process on $T_H^{(d)}$ with root sequence X and let Y_1, \dots, Y_n be the sequences obtained at the leaves, where $Y_i = y_1^i, \dots, y_{k_i}^i$.

Theorem 1.2 (Main result) *For all $\chi > 0$, there is $\Phi, \Phi', \Phi'' > 0$ and $d'' > 0$ such that the following holds for $d \geq d''$ and $\beta = d^{-1}$. There is a polynomial-time algorithm \mathbb{A} with access to Y_1, \dots, Y_n such that for all*

$$(1 - 2p_s)^2 > \frac{\Phi \log d}{d},$$

$$p_i + p_d < \frac{\Phi'}{k^{2/3} \log n},$$

$$\Phi'' \log^3 n < k < \text{poly}(n),$$

the algorithm \mathbb{A} outputs a binary sequence \widehat{X} which satisfies the following with probability at least $1 - \chi$:

1. $\widehat{X} = \widehat{x}_1, \dots, \widehat{x}_k$ has length k .
2. For all $j = 1, \dots, k$, $\mathbb{P}[\widehat{x}_j = x_j] > 1 - \beta$.

Proof sketch. We give a brief proof sketch. As discussed previously, in the presence of indels the reconstruction of ancestral sequences requires the solution of the *multiple sequence alignment* problem. However, in addition to being computationally intractable, global alignment through the optimization of a pairwise scoring function may create biases and correlations that are hard to quantify. Therefore, we require a more probabilistic approach. From a purely information-theoretic point of view the pairwise alignment of sequences that are far apart in the tree is difficult. A natural solution to this problem is instead to perform *local* alignments and ancestral reconstructions, and recurse our way up the tree.

This *recursive* approach raises its own set of issues. Consider a parent node and its d children. It may be easy to perform a local alignment of the children’s sequences and derive a good approximation to the parent sequence (for example, through site-wise majority). Note however that, to allow a recursion of this procedure all the way to the root, we have to provide strong guarantees about the probabilistic behavior of our local ancestral reconstruction. As is the case for global alignment, a careless alignment procedure creates biases and correlations that are hard to control. For instance, it is tempting to treat misaligned sites as independent unbiased noise but this idea presents difficulties:

Consider a site j of the parent sequence and suppose that for this site we have succeeded in aligning all but two of the children, say 1 and 2. Let $x_{j_i}^i$ denote the site in the i ’th child which was used to estimate the j ’th site. By the independence assumption on the root sequence and the inserted sites, $x_{j_1}^1$ and $x_{j_2}^2$ are uniform and independent of $(x_{j_i}^i)_{i=3}^d$. However, $x_{j_1}^1$ and $x_{j_2}^2$ may originate from the *same* neighboring site of the parent sequence and therefore are themselves correlated.

Quantifying the effect of this type of correlation appears to be nontrivial.

Instead, we use an *adversarial* approach to local ancestral reconstruction. That is, we treat the misaligned sites as being controlled by an adversary who seeks to flip the reconstructed value. This comes at a cost: it produces an asymmetry in our ancestral

reconstruction. Although the RP is well-studied in the symmetric noise case, much remains to be understood in the asymmetric case. In particular, obtaining tight results in terms of substitution probability here may not be possible as the critical threshold of the RP may be hard to identify. We do however provide a tailored analysis of the particular instance of the RP by recursive majority obtained through this adversarial approach and we obtain results that are close to the known threshold for the symmetric case. Unlike the standard RP, the reconstruction error is not i.i.d. but we show instead that it “dominates” an i.i.d. noise. (See Section 4.2 for a definition.) This turns out to be enough for a well-controlled recursion. We first define a local alignment procedure which has a fair success probability (independent of n). However, applying this alignment procedure multiple times in the tree is bound to fail sometimes. We therefore prove that the local reconstruction procedure is somewhat robust in the sense that even if one of the d inputs to the reconstruction procedure is faulty, it still has a good probability of success.

As for our local alignment procedure, we adopt an *anchor* approach. Anchors were also used by [4, 5]—although in a quite different way. We imagine a partition of every node’s sequence into islands of length $O(k^{1/3})$. (The precise choice of the island length comes from a trade-off between the length and the number of islands in bounding the “bad” events below.) At the beginning of each island we have an anchor of length $O(\log n)$. Through this partition of the sequences in islands and anchors we aim to guarantee the following. Given a specific father node v , with fair probability 1) all the anchors in the children nodes are indel-free; and 2) for all parent islands, almost all of the corresponding children islands have no indel at all and, moreover, at most one child island may have a single indel. The “bad” children islands—those that do not satisfy these properties—are treated as controlled by an adversary. We show that Conditions 1) and 2) are sufficient to guarantee that: the anchors of all islands can be aligned with high probability and single indel events between anchors can be identified. This allows a local alignment of all islands with at most one “bad” child per island and is enough to perform a successful adversarial recursive majority vote as described above. The bound on the maximum indel probability sustained by our reconstruction algorithm comes from satisfying Conditions 1) and 2) above.

Notation. For a sequence $X = x_1, \dots, x_k$, we let $X[i : j] = x_i, \dots, x_j$. We use the expression “with high probability (w.h.p.)” to mean “with probability at least $1 - 1/\text{poly}(n)$ ” where the polynomial in n

can be made of arbitrarily high degree (by choosing the appropriate constants large enough). We denote by $\text{Bin}(n, p)$ a random variable with binomial distribution of parameters n, p . For two random variables X, Y we denote by $X \sim Y$ the equality in distribution.

Organization. The rest of the paper is organized as follows. We describe the algorithm in Section 2. The proof of our main result is divided into two sections. In Section 3, we state a series of high-probability claims about the evolutionary process. Then, conditioning on these claims, we provide a deterministic analysis of the correctness of the algorithm in Section 5.2. All proofs are omitted from this extended abstract.

2 Description of the Algorithm

In this section we describe our algorithm for TRPT. Our algorithm is recursive, proceeding from the leaves of the tree to the root. We describe the recursive step applied to a non-leaf node of the tree.

Recursive Setup—Our Goal. For our discussion in this section, let us consider a non-leaf node v with d children, denoted u_i for $i \in [d]$. For notational convenience, we drop the index u and denote its children by $1, \dots, d$. Our goal for the recursive step of the algorithm is to reconstruct the sequence at the node v given the sequences of the children. Denote the sites of the father by $X_0 = x_1^0, \dots, x_{k_0}^0$, and the sites of the i 'th child by $X_i = x_1^i, \dots, x_{k_i}^i$. During the reconstruction process, we do not have access to the children's sequences, but rather to reconstructed sequences denoted by $\hat{X}_i = \hat{x}_1^i, \dots, \hat{x}_{k_i}^i$.

Let us consider the following partition of the sequence of v into subsequences, called *islands*. Of course our algorithm doesn't have access to the sequence at v during the recursive step of the algorithm. We define the partition as a means to describe our algorithm: The sites of v are partitioned into *islands* of length $\ell = k^{1/3}$ (except for the last one which is possibly shorter). Denote by $N_0 = \lceil k_0/\ell \rceil$ the number of islands in v . Each island starts with an *anchor* of a bits. That is, the islands are the bitstrings $X_0[1 : \ell], X_0[\ell + 1 : 2\ell], \dots$ and the anchors are the bitstrings $X_0[1 : a], X_0[\ell + 1 : \ell + a], \dots$

Our algorithm tries to identify for each island $X_0[(i - 1)\ell + 1 : i\ell]$ the substrings of each of the d children that correspond to this island (i.e., contain the sites of the island), called "child islands." We do so iteratively for $i = 1 \dots N_0$. We use the islands that did not have indels for sequence reconstruction, using the substitution-only model. Some islands will have indels however. This leads to two "modes of failure":

one invalidates the entire (parent) node, and the other invalidates only an island of a child. More specifically, a node becomes invalidated (i.e., useless) when indels are not evenly distributed, that is: when an indel occurred in an anchor, or two (or more) indels occurred in a specific island over all d children. This is a rare event. Barring this event, we expect that each island suffers only at most one indel over all children. The island (of a child) that has exactly one indel is invalidated (second mode of failure), and is thus deemed useless for reconstruction purposes. As long as the parent node is not invalidated, each island will have at least $d - 2$ non-invalidated children islands (one additional island is potentially lost to a child node that may have been invalidated at an earlier stage).

Even when the algorithm identifies that a child island has an indel somewhere, the island is not ignored. The algorithm still needs to compute the length of the island in order to know the start of the next island in this child. For this purpose, we use the anchor of the next island and match it to the corresponding anchors of the other (non-invalidated) child islands. In fact the same procedure lets us detect which of the child islands are invalidated.

More formally, we define d functions $f_i : \{1, \dots, k_0\} \rightarrow \{1, \dots, k_i\} \cup \{\dagger\}$, where f_i takes the sites of v to the corresponding sites of the i 'th child or to the special symbol \dagger if the site was deleted. Note that for each i , f_i is monotone, when ignoring sites which are mapped to \dagger . For $t = \ell r$, let $s_i(r) = f_i(t+1) - (t+1)$ denote the displacement of the site corresponding to the $(t + 1)$ 'st site of the parent, in the i 'th child. By convention, we take $s_i(0) = 0$. If there is no indel between $t = \ell r$ and $t' = \ell r'$ then $s_i(r) = s_i(r')$. Note that, in the specific case of one indel operation in the island, we have that $|s_i(r) - s_i(r')| = 1$.

Algorithm. Our algorithm estimates the values of $s_i(r)$ and uses these estimates to match the starting positions of the islands in the children. The full algorithm is given in Figure 1. We use the following additional notation. For $x \in \{0, 1\}$ we let $\langle x \rangle = 2x - 1$. Then, for two $\{0, 1\}^m$ -sequences $Y = y_1, \dots, y_m$ and $Z = z_1, \dots, z_m$, we define their (empirical) correlation as

$$\text{Corr}(Y, Z) = \frac{1}{m} \sum_{j=1}^m \langle y_j \rangle \langle z_j \rangle.$$

Note that $y \mapsto \langle y \rangle$ maps 1 to 1 and 0 to -1 . One can think of $\text{Corr}(Y, Z)$ as a form of normalized centered Hamming distance between Y and Z . In particular, a large value of $\text{Corr}(Y, Z)$ implies that Y and Z tend to agree. We will use the following threshold (which




Figure 1: This is the basic recursive step of our reconstruction algorithm. It takes as input the d inferred sequences of the children $\hat{x}^1, \dots, \hat{x}^d$ and computes a sequence for the parent \hat{x}^0 . If any of the steps above cannot be accomplished, we abort the reconstruction of the parent and declare it radioactive.

will be justified in Section 5.1)

$$\gamma = ((1 - \delta)(1 - 2p_s)^2 - 4\beta),$$

where δ is chosen so that

$$(1 - \delta)(1 - 2p_s)^2 - 8\beta > \delta + 8\beta,$$

where again $\beta = d^{-1}$ and d is large enough.

3 Analyzing the Indel Process

We define $a \geq C \log n$ and $\alpha \leq \varepsilon/d < 1$, for constants C, ε to be determined later. We require $a < k^{1/3} < \text{poly}(n)$. We assume that the indel probability per site satisfies

$$p_{\text{id}} = \frac{\alpha}{4dk^{2/3}a} = O\left(\frac{1}{k^{2/3} \log n}\right).$$

Throughout, we denote the tree by $T = (V, E)$.

3.1 Bound on the Sequence Length

As the indel probability is defined per site, longer sequences suffer more indel operations than shorter ones. We begin by bounding the effect of this process. We claim that with high probability the lengths of all sequences are roughly equal.

Claim 3.1 (Bound on sequence length) *For all $\zeta > 0$ (small), there exists $C' > 0$ (large) so that for all u in V , we have*

$$k_v \in [\underline{k}, \bar{k}] \equiv [(1 - \zeta)k, (1 + \zeta)k],$$

with high probability given $k \geq C' \log^3 n$. We denote this event by \mathcal{L} .

3.2 Existence of a Dense Stable Subtree

We claim that with probability close to 1 there exists a dense subtree of T with a “good indel structure,” as defined below. Our algorithm will try to identify this subtree and perform reconstruction on it, as described in Section 4.

Indel structure of a node. Recall that $\ell = k^{1/3}$.

Definition 3.2 (Indel structure) *For a node (parent) v , we say that v is radioactive if one of the following events happen:*

1. *Event \mathcal{B}_1 : Node v has a child u such that when evolving from v to u an indel operation occurred in at least one of the sites which are located in an anchor.*
2. *Event \mathcal{B}_2 : There is an island I and two children u, u' , such that an indel occurred in I in the transition from v to u and in the transition from v to u' .*
3. *Event \mathcal{B}_3 : There is an island I and a child u , such that two indel operations (or more) happened in I in the transition from v to u .*

Otherwise the node v is stable. By definition, the leaves of T are stable. A subtree of T is stable if all of its nodes are stable.

Claim 3.3 (Bound on radioactivity) *For all $\alpha > 0$, there exists a choice of $\zeta > 0$ small enough in Claim 3.1 such that conditioning on the event \mathcal{L} occurring: any vertex v is radioactive with probability at most α .*

Claim 3.4 (Existence of a dense stable subtree) *For all $\chi > 0$, there is a choice of $\zeta > 0$ small enough in Claim 3.1 such that, conditioning on the event \mathcal{L} occurring, with probability at least $1 - \chi$, the root of T is the father of a $(d - 1)$ -ary stable subtree of T . We denote this event by \mathcal{S} .*

4 A Stylized Reconstruction Process

We describe a hypothetical sequence reconstruction process performed on the stable tree defined by the indels. We argue that the process gives strong reconstruction guarantees (assuming that the radioactive nodes and the islands with indels are controlled by an adversary). Then we argue that our algorithm performs at least as well as the reconstruction process against the adversary. Throughout we suppose that a stable tree exists and is given to us, together with the “orbit” of every site of the sequence at the root of the tree (see function F below). However, we are given no information about the substitution process.

Let $v \in V$ and assume v is the root of a $(d - 1)$ -ary stable subtree $T^* = (V^*, E^*)$ of T . (We make

the stable subtree below v into a $(d - 1)$ -ary tree by potentially removing arbitrary nodes from it, at random.) Let $u \in V^*$. For each island I in u , at most one child u' of u in T^* contains an indel in which case it contains exactly one indel. We say that such an I is a corrupted island of u' . The basic intuition behind our analysis is that, provided the alignment on T^* is performed correctly (which we defer to Section 5.2), the ancestral reconstruction step of our algorithm is a recursive majority procedure against an adversary which controls the corrupted islands and the radioactive nodes (as well as all their descendants). Below we analyze this adversarial process.

Recursive majority. We begin with a formal definition of recursive majority. Let $\text{Maj} : \{0, 1, \#\}^d \rightarrow \{0, 1\}$ be the function that returns the majority value over non- $\#$ values, and flips an unbiased coin in case of a tie (including the all- $\#$ vector). Let $n_0 = d^{H_0}$ be the number of leaves in T below v . Consider the following recursive function of $z = (z_1, z_2, \dots, z_{n_0}) \in \{0, 1, \#\}^{n_0}$: $\text{Maj}^0(z_1) = z_1$, and

$$\begin{aligned} \text{Maj}^j(z_1, \dots, z_{d^j}) \\ &= \text{Maj}(\text{Maj}^{j-1}(z_1, \dots, z_{d^{j-1}}), \dots, \\ &\quad \text{Maj}^{j-1}(z_{d^j-d^{j-1}+1}, \dots, z_{d^j})), \end{aligned}$$

for all $j = 1, \dots, H_0$. Then, $\text{Maj}^{H_0}(z)$ is the d -wise recursive majority of z .

Let $X_0 = x_1^0, \dots, x_{k_0}^0$ be the sequence at v . For $u \in V^*$ and $t = 1, \dots, k_0$, we denote by $F_u(t)$ the position of site x_t^0 in u or \dagger if the site has been deleted on the path to u . We say that $\mathcal{C}_{u,t}$ holds if $F_u(t)$ is in a corrupted island of u . Let $\text{Path}(u, v)$ be the set of nodes on the path between u and v .

Definition 4.1 (Gateway node) *A node u is a gateway for site t if:*

1. $F_u(t) \neq \dagger$; and
2. For all $u' \in \text{Path}(u, v) - \{v\}$, $\mathcal{C}_{u',t}$ does not hold.

We let $T_t^{**} = (V_t^{**}, E_t^{**})$ be the subtree of T^* containing all gateway nodes for t . By construction, T_t^{**} is at least $(d - 2)$ -ary and for convenience we remove arbitrary nodes, at random, to make it exactly $(d - 2)$ -ary. Notice that, for $t, t' \in [1 : k_0]$, the subtrees T_t^{**} and $T_{t'}^{**}$ are random and correlated. However, they are independent of the substitution process.

We will argue in Section 5.2 that the reconstructed sequence produced by our method at v “dominates” (see below) the following reconstruction process. Let $L_v = u_1, \dots, u_{n_0}$ be the leaves below v ordered according to a planar realization of the subtree below v . Denote by $X_i = x_1^i, \dots, x_{k_i}^i$ the sequence at u_i . For $t = 1, \dots, k_0$, let L_t^{**} be the leaves of T_t^{**} . We define

the following auxiliary sequences: for $u_i \in L_v$, we let $\Xi_i = \xi_1^i, \dots, \xi_{k_i}^i$ where for $t = 1, \dots, k_0$

$$\xi_t^i = \begin{cases} x_{F_{u_i}(t)}^i & \text{if } u_i \in L_t^{**} \\ 1 - x_t^0 & \text{o.w.} \end{cases}$$

In words, ξ_t^i is the descendant of x_t^0 if u_i is a gateway to t and is the opposite of the value x_t^0 otherwise. Because of the monotonicity of recursive majority, the latter choice is in some sense the “worst adversary” (ignoring correlations between sites—we will come back to this point later). We then define a reconstructed sequence at v as $\widehat{\Xi}_0 = \widehat{\xi}_1^0, \dots, \widehat{\xi}_{k_0}^0$ where for $t = 1, \dots, k_0$

$$\widehat{\xi}_t^0 = \text{Maj}^{H_0}(\xi_t^1, \dots, \xi_t^{n_0}).$$

We now analyze the accuracy of this (hypothetical) estimator—which we refer to as the *adversarial reconstruction* of X_0 . We argue in Section 5.2 that our actual estimator is at least as good as $\widehat{\Xi}_0$ w.h.p.

4.1 Recursive Majority Against an Adversary

To analyze the performance of the adversarial reconstruction $\widehat{\Xi}_0$, we consider the following stylized process.

Definition 4.2 (Adversarial Process) We consider the following process:

1. Run the evolutionary process on $T_{H_0}^{(d-2)}$ at one position only starting with root state 0 without indels, that is, taking $p_{\text{id}} = 0$.
2. Then complete $T_{H_0}^{(d-2)}$ into $T_{H_0}^{(d)}$ and associate to each additional node the state 1.
3. Let $R_{H_0}^{(d)}$ be the random variable in $\{0, 1\}$ obtained by running recursive majority on the leaf states obtained above.

We call this process the *recursive majority against an adversary* on $T_{H_0}^{(d)}$.

Claim 4.3 (Accuracy of recursive majority)

There exists a constant $C'' > 0$ and $d'' > 0$ such that taking

$$\theta_s^2 > \frac{C'' \log d}{d},$$

and $d \geq d''$, then the probability that the recursive majority against an adversary on $T_{H_0}^{(d)}$ correctly reconstructs root state 0 is at least $1 - \beta$ uniformly in H_0 where $\beta = d^{-1}$. In comparison, note that the Kesten-Stigum bound for binary symmetric channels on d -ary trees is $\theta^2 > d^{-1}$ [32, 54].

Definition 4.4 (Bernoulli sequence) For $q > 0$ and $m \in \mathbb{N}$, the (q, m) -Bernoulli sequence is the product distribution on $\{0, 1\}^m$ such that each position is 1 independently with probability $1 - q$. We denote by $B_{q,m}$ the corresponding random variable.

Claim 4.5 (Subsequence reconstruction)

Assume v is the root of a $(d - 1)$ -ary stable subtree. Choosing $C'' > 0$ and $d'' > 0$ as in Claim 4.3 is such that the following holds for $d \geq d''$ and $\beta = d^{-1}$. For $t, m \in \{1, \dots, k_0\}$, let $\Lambda = (\lambda_1, \dots, \lambda_m)$ be the agreement vector between the $\widehat{\Xi}_0[t + 1 : t + m]$ and $X_0[t + 1 : t + m]$, that is, $\lambda_i = 1$ if recursive majority correctly reconstructs position i . Then there is $0 \leq \beta' \leq \beta$ such that $\Lambda \sim B_{\beta', m}$. (Here, β' may depend on H_0 but β does not.)

4.2 Stochastic Domination and Correlation

In our discussion so far we have assumed that a stable tree exists and is given to us, together with the function F . This allowed us to define the stylized recursive majority process against an adversary for which we claimed strong reconstruction guarantees. In reality, we have no access to the stable tree. We construct it recursively from the leaves towards the root. At the same time we align sequences, discover corrupted islands, and reconstruct sequences of internal nodes. The stylized recursive majority process may be used to provide a “lower bound” on the actual reconstruction process. The notion of “lower bound” that is of interest to us is captured by *stochastic domination*, which we proceed to define formally.

Definition 4.6 (Stochastic domination) Let X, Y be two random variables in $\{0, 1\}^m$. We say that Y stochastically dominates X , denoted $X \preceq Y$, if there is a joint random variable (\tilde{X}, \tilde{Y}) such that the marginals satisfy $X \sim \tilde{X}$ and $Y \sim \tilde{Y}$ and moreover $\mathbb{P}[\tilde{X} \leq \tilde{Y}] = 1$.

Note that in the definition above X and Y may (typically) live in different probability spaces. Then, the joint variable (\tilde{X}, \tilde{Y}) is a coupled version of X and Y . In our case, X is the adversarial recursive process whereas Y is the actual reconstruction performed by the algorithm. We now explain how to use this property for correlation estimation.

Correlation. The analysis of the previous section guarantees that the sequences output by the adversarial reconstruction process are well correlated with the true sequences. But if we are only going to use the adversarial process as a lower bound for the true reconstruction process, it is important to establish that stochastic domination preserves correlation. In preparing the ground for such a claim

let us establish an important property of the adversarial process. Let T_u and T_v be the two disjoint copies of $T_h^{(d)}$ rooted at the nodes u and v respectively, and let $X = x_1, x_2, \dots, x_m \in \{0, 1\}^m$ and $Y = y_1, y_2, \dots, y_m \in \{0, 1\}^m$ be sequences at the nodes u and v . Assume that u and v are the roots of $(d-1)$ -ary stable subtrees. Let $\widehat{X}' = \widehat{x}'_1, \widehat{x}'_2, \dots, \widehat{x}'_m \in \{0, 1\}^m$ and $\widehat{Y}' = \widehat{y}'_1, \widehat{y}'_2, \dots, \widehat{y}'_m \in \{0, 1\}^m$ be the reconstructions of X and Y obtained by the adversarial reconstruction process. Let $\Lambda = \lambda_1, \dots, \lambda_m$ and $\Theta = \theta_1, \dots, \theta_m$ be the resulting agreement vectors.

Claim 4.7 (Concentration of bias) *Let β', β be as in Claim 4.5. Then, with probability at least $1 - e^{-\Omega(m\beta^2)}$ the following are satisfied*

$$\left| \frac{1}{m} \sum_{i=1}^m \langle \lambda_i \rangle \langle \theta_i \rangle - (1 - 2\beta')^2 \right| \leq \frac{1}{2}\beta;$$

$$\left| \frac{1}{m} \sum_{i=1}^m \mathbb{1}_{\langle \lambda_i \rangle = -1} - \beta' \right| \leq \frac{1}{2}\beta;$$

$$\left| \frac{1}{m} \sum_{i=1}^m \mathbb{1}_{\langle \theta_i \rangle = -1} - \beta' \right| \leq \frac{1}{2}\beta.$$

Claim 4.8 (Correlation bound) *Let $\widehat{X}, \widehat{Y} \in \{0, 1\}^m$ be random strings defined on the same probability space as \widehat{X}' and \widehat{Y}' . Denote by Z (resp. W) the agreement vectors of \widehat{X} (resp. \widehat{Y}) with X (resp. Y). Assume that $\Lambda \leq Z$ and $\Theta \leq W$ with probability 1, where Λ and Θ are the agreement vectors of \widehat{X}' and \widehat{Y}' with X and Y as explained above. Then, conditioned on the conclusions of Claim 4.7, we have, with probability 1*

$$|\text{Corr}(X, Y) - \text{Corr}(\widehat{X}, \widehat{Y})| \leq 8\beta.$$

5 Analyzing the True Reconstruction Process

In Section 5.1, we argue that, if a stable subtree exists, the adversarial reconstructions of aligned anchors exhibit strong correlation signal, while misaligned anchors exhibit weak signal. This holds true for sequences that stochastically dominate the adversarial reconstructions.

5.1 Anchor Alignment

Consider a parent v that is stable. Let i, j be two children with sequences $X_i = x_1^i, \dots, x_{k_i}^i$ and $X_j = x_1^j, \dots, x_{k_j}^j$. Let $t = \ell r$ and consider the following subsequences (of length a) at i and j

$$\mathcal{A}_r^i = x^i[t + s_i(r) + 1 : t + s_i(r) + a],$$

and

$$\mathcal{A}_r^j = x^j[t + s_j(r) + 1 : t + s_j(r) + a].$$

These are related (but not identical) to the definition of anchors in the algorithm of Section 2. In particular, note that by definition \mathcal{A}_r^i and \mathcal{A}_r^j are always aligned, in the sense that they correspond to the same subsequence of v . Consider also the following subsequences

$$\mathcal{D}_r^j = x^j[t + s_j(r) : t + s_j(r) + a - 1],$$

and

$$\mathcal{I}_r^j = x^j[t + s_j(r) + 2 : t + s_j(r) + a + 1].$$

These are the one-site shifted subsequences for j . We claim that \mathcal{A}_r^i is always significantly more correlated to its aligned brother \mathcal{A}_r^j than to the misaligned ones \mathcal{D}_r^j and \mathcal{I}_r^j . This follows from the fact that the misaligned subsequences are sitewise independent.

Claim 5.1 (Anchor correlations) *For all $\delta > 0$ (and d large enough) such that $(1 - \delta)(1 - 2p_s)^2 - 8\beta > \delta + 8\beta$, there is $C > 0$ large enough so that with $a = C \log n$, the following hold:*

1. **Aligned anchors.**

$$\begin{aligned} \mathbb{P}[\text{Corr}(\mathcal{A}_r^i, \mathcal{A}_r^j) > (1 - \delta)(1 - 2p_s)^2] \\ > 1 - \exp(-\Omega(a)) = 1 - 1/\text{poly}(n). \end{aligned}$$

2. **Misaligned anchors.**

$$\begin{aligned} \mathbb{P}[\text{Corr}(\mathcal{A}_r^i, \mathcal{D}_r^j) < \delta] \\ > 1 - \exp(-\Omega(a)) = 1 - 1/\text{poly}(n), \end{aligned}$$

and similarly for \mathcal{I}_r^j .

We denote by $\mathcal{A}_{i,j,r}$ the above events and their symmetric counterparts under $i \leftrightarrow j$.

Claim 5.2 (Reconstructed version) *Let $\widehat{X}_i = (\widehat{x}_i^i)_{i=1}^{k_i}$ and $\widehat{X}_j = (\widehat{x}_i^j)_{i=1}^{k_j}$ dominate the adversarial reconstructions \widehat{X}'_i and \widehat{X}'_j of X_i and X_j , as defined in Claim 4.8. Let $\widehat{\mathcal{A}}_r^i = \widehat{x}^i[t + s_i(r) + 1 : t + s_i(r) + a]$ and similarly for all other possibilities $\widehat{\mathcal{A}} \leftrightarrow \widehat{\mathcal{D}}, \widehat{\mathcal{I}}$ and/or $i \leftrightarrow j$. Denote by $\mathcal{B}_{i,j,r}$ the event that the conclusions of Claim 4.7 hold for \widehat{X}'_i and \widehat{X}'_j over all pairs of intervals involving $[t + s_i(r) : t + s_i(r) + a - 1]$, $[t + s_i(r) + 1 : t + s_i(r) + a]$, and $[t + s_i(r) + 2 : t + s_i(r) + a + 1]$, with $i \leftrightarrow j$ as necessary. Then, conditioned on $\mathcal{B}_{i,j,r}$ and $\mathcal{A}_{i,j,r}$ we have*

$$\text{Corr}(\widehat{\mathcal{A}}_r^i, \widehat{\mathcal{A}}_r^j) > (1 - \delta)(1 - 2p_s)^2 - 8\beta,$$

$$\text{Corr}(\widehat{\mathcal{A}}_r^i, \widehat{\mathcal{D}}_r^j) < \delta + 8\beta,$$

and

$$\text{Corr}(\widehat{\mathcal{A}}_r^i, \widehat{\mathcal{I}}_r^j) < \delta + 8\beta,$$

as well as their symmetric counterparts under $i \leftrightarrow j$.

5.2 Proof of Correctness

We claim that our recursive procedure reconstructs the desired sequence at the root of the tree whenever a collection of good events occurs. Recall the definitions of the events \mathcal{L} , \mathcal{S} , $\mathcal{B}_{i,j,r}$, $\mathcal{A}_{i,j,r}$ from Claims 3.1, 3.4, 5.1 and 5.2. Conditioning on \mathcal{L} and \mathcal{S} , denote by $T^* = (V^*, E^*)$ the stable $(d-1)$ -ary subtree of T . Then, for all $v \in V^*$, all pairs of children i, j of v in T^* , and all $r = 1, \dots, \bar{k}/\ell$, we condition on the events $\mathcal{B}_{i,j,r}$ and $\mathcal{A}_{i,j,r}$. Note that having conditioned on \mathcal{L} there is only a polynomial number of such events, since all sequence lengths are bounded by \bar{k} . (If $r \cdot \ell$ is larger than a node's sequence length we assume that the corresponding events are vacuously satisfied.) Finally recall that, conditioning on \mathcal{L} , the event \mathcal{S} occurs with probability $1 - \chi$ and all other events occur with high probability. We denote the collection of events by \mathcal{E} .

Conditioning on \mathcal{E} , the proof of correctness of the algorithm follows from a bottom-up induction. The gist of the argument is the following. Suppose that at a recursive step of the algorithm we have reconstructed sequences for all children of a node v , which are strongly correlated with the true sequences (in the sense of dominating the corresponding adversarial reconstructions). Having conditioned on the events $\mathcal{A}_{i,j,r}$ and $\mathcal{B}_{i,j,r}$, it follows then that the correct alignments of anchors exhibit strong correlation signal while the incorrect alignments weak correlation signal. Hence, our correlation tests between anchors discover the corrupted islands and do the anchor alignments correctly (at least for all nodes lying inside the stable tree). Hence the shift functions \hat{s}_i 's are correctly inferred, and the reconstruction of v 's sequence can be shown to dominate the corresponding adversarial reconstruction.

Acknowledgments

C.D. A.H. and S.R. performed part of this work at Microsoft Research. A.H. thanks the W. M. Keck Foundation Center for Extreme Quantum Information Theory for their support.

References

- [1] Vladimir I. Levenshtein. Efficient reconstruction of sequences. *IEEE Transactions on Information Theory*, 47(1):2–22, 2001.
- [2] Vladimir I. Levenshtein. Efficient reconstruction of sequences from their subsequences or supersequences. *J. Comb. Theory Ser. A*, 93(2):310–332, 2001. ISSN 0097-3165. doi: <http://dx.doi.org/10.1006/jcta.2000.3081>.
- [3] Tuğkan Batu, Sampath Kannan, Sanjeev Khanna, and Andrew McGregor. Reconstructing strings from random traces. In *SODA '04: Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 910–918, Philadelphia, PA, USA, 2004. Society for Industrial and Applied Mathematics. ISBN 0-89871-558-X.
- [4] S. Kannan and A. McGregor. More on reconstructing strings from random traces: insertions and deletions. In *Proceedings of ISIT*, pages 297–301, 2005.
- [5] Thomas Holenstein, Michael Mitzenmacher, Rina Panigrahy, and Udi Wieder. Trace reconstruction with constant deletion probability and related results. In *SODA '08: Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 389–398, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.
- [6] Krishnamurthy Viswanathan and Ram Swaminathan. Improved string reconstruction over insertion-deletion channels. In *SODA '08: Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 399–408, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.
- [7] J. W. Thornton. Resurrecting ancient genes: experimental analysis of extinct molecules. *Nat. Rev. Genet.*, 5(5):366–375, 2004.
- [8] C. Semple and M. Steel. *Phylogenetics*, volume 22 of *Mathematics and its Applications series*. Oxford University Press, 2003.
- [9] J. A. Cavender. Taxonomy with confidence. *Math. Biosci.*, 40(3-4), 1978.
- [10] J. S. Farris. A probability model for inferring evolutionary trees. *Syst. Zool.*, 22(4):250–256, 1973.
- [11] J. Neyman. Molecular studies of evolution: a source of novel statistical problems. In S. S. Gupta and J. Yackel, editors, *Statistical decision theory and related topics*, pages 1–27. Academic Press, New York, 1971.
- [12] E. Mossel. On the impossibility of reconstructing ancestral data and phylogenies. *J. Comput. Biol.*, 10(5):669–678, 2003. URL http://www.stat.berkeley.edu/~mossel/publications/jcb_impossibility.pdf.
- [13] E. Mossel. Phase transitions in phylogeny. *Trans. Amer. Math. Soc.*, 356(6):2379–2404, 2004. ISSN 0002-9947. URL http://www.stat.berkeley.edu/~mossel/publications/phylogeny_ising.ps.
- [14] Constantinos Daskalakis, Elchanan Mossel, and Sébastien Roch. Optimal phylogenetic reconstruction. In *STOC'06: Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 159–168, New York, 2006. ACM.
- [15] Sébastien Roch. Sequence-length requirement for distance-based phylogeny reconstruction: Breaking the polynomial barrier. In *FOCS*, pages 729–738, 2008.
- [16] E. Mossel. Recursive reconstruction on periodic trees. *Random Struct. Algor.*, 13(1):81–97, 1998.

- ISSN 1042-9832. URL <http://www.stat.berkeley.edu/~mossel/publications/recursive.ps>.
- [17] W. S. Evans, C. Kenyon, Y. Peres, and L. J. Schulman. Broadcasting on trees and the Ising model. *Ann. Appl. Probab.*, 10(2):410–433, 2000.
- [18] E. Mossel. Reconstruction on trees: beating the second eigenvalue. *Ann. Appl. Probab.*, 11(1):285–300, 2001. ISSN 1050-5164. URL <http://www.stat.berkeley.edu/~mossel/publications/second.ps>.
- [19] E. Mossel and Y. Peres. Information flow on trees. *Ann. Appl. Probab.*, 13(3):817–844, 2003. ISSN 1050-5164. URL http://www.stat.berkeley.edu/~mossel/publications/inf_flow.ps.
- [20] F. Martinelli, A. Sinclair, and D. Weitz. Glauber dynamics on trees: boundary conditions and mixing time. *Comm. Math. Phys.*, 250(2):301–334, 2004.
- [21] S. Janson and E. Mossel. Robust reconstruction on trees is determined by the second eigenvalue. *Ann. Probab.*, 32:2630–2649, 2004. URL <http://www.stat.berkeley.edu/~mossel/publications/robust.pdf>.
- [22] N. Berger, C. Kenyon, E. Mossel, and Y. Peres. Glauber dynamics on trees and hyperbolic graphs. *Probab. Theory Rel.*, 131(3):311–340, 2005. Extended abstract by Kenyon, Mossel and Peres appeared in proceedings of 42nd IEEE Symposium on Foundations of Computer Science (FOCS) 2001, 568–578.
- [23] Christian Borgs, Jennifer T. Chayes, Elchanan Mossel, and Sébastien Roch. The Kesten-Stigum reconstruction bound is tight for roughly symmetric binary channels. In *FOCS*, pages 518–530, 2006.
- [24] Antoine Gerschenfeld and Andrea Montanari. Reconstruction for models on random graphs. In *FOCS '07: Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, pages 194–204, Washington, DC, USA, 2007. IEEE Computer Society. ISBN 0-7695-3010-9. doi: <http://dx.doi.org/10.1109/FOCS.2007.58>.
- [25] N. Bhatnagar, J. Vera, E. Vigoda, and Weitz D. Reconstruction for colorings on trees. Preprint available at arxiv.org/abs/0711.3664, 2007.
- [26] A. Sly. Reconstruction of random colourings. *Communications in Mathematical Physics*, 288(3):943–961, 2009.
- [27] Allan Sly. Reconstruction for the Potts model. In *STOC '09: Proceedings of the 41st annual ACM symposium on Theory of computing*, pages 581–590, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-506-2. doi: <http://doi.acm.org/10.1145/1536414.1536493>.
- [28] S. Roch. *Markov Models on Trees: Reconstruction and Applications*. PhD thesis, UC Berkeley, 2007.
- [29] A. Sly. *Spatial and Temporal Mixing of Gibbs Measures*. PhD thesis, UC Berkeley, 2009.
- [30] P. M. Bleher, J. Ruiz, and V. A. Zagrebnov. On the purity of the limiting Gibbs state for the Ising model on the Bethe lattice. *J. Statist. Phys.*, 79(1-2):473–482, 1995.
- [31] D. Ioffe. On the extremality of the disordered state for the Ising model on the Bethe lattice. *Lett. Math. Phys.*, 37(2):137–143, 1996.
- [32] H. Kesten and B. P. Stigum. Additional limit theorems for indecomposable multidimensional Galton-Watson processes. *Ann. Math. Statist.*, 37:1463–1481, 1966.
- [33] Ari Loytynoja and Nick Goldman. Phylogeny-Aware Gap Placement Prevents Errors in Sequence Alignment and Evolutionary Analysis. *Science*, 320(5883):1632–1635, 2008. doi: 10.1126/science.1158395. URL <http://www.sciencemag.org/cgi/content/abstract/320/5883/1632>.
- [34] Karen M. Wong, Marc A. Suchard, and John P. Huelsenbeck. Alignment Uncertainty and Genomic Analysis. *Science*, 319(5862):473–476, 2008. doi: 10.1126/science.1151532. URL <http://www.sciencemag.org/cgi/content/abstract/319/5862/473>.
- [35] Jeffrey L. Thorne, Hirohisa Kishino, and Joseph Felsenstein. An evolutionary model for maximum likelihood alignment of dna sequences. *Journal of Molecular Evolution*, 33(2):114–124, 1991.
- [36] Jeffrey L. Thorne, Hirohisa Kishino, and Joseph Felsenstein. Inching toward reality: An improved likelihood model of sequence evolution. *Journal of Molecular Evolution*, 34(1):3–16, 1992.
- [37] Dirk Metzler. Statistical alignment based on fragment insertion and deletion models. *Bioinformatics*, 19(4):490–499, 2003. doi: 10.1093/bioinformatics/btg026. URL <http://bioinformatics.oxfordjournals.org/cgi/content/abstract/19/4/490>.
- [38] I. Miklos, G. A. Lunter, and I. Holmes. A "Long Indel" Model For Evolutionary Sequence Alignment. *Mol Biol Evol*, 21(3):529–540, 2004. doi: 10.1093/molbev/msh043. URL <http://mbe.oxfordjournals.org/cgi/content/abstract/21/3/529>.
- [39] Marc A. Suchard and Benjamin D. Redelings. BALi-Phy: simultaneous Bayesian inference of alignment and phylogeny. *Bioinformatics*, 22(16):2047–2048, 2006. doi: 10.1093/bioinformatics/btl175. URL <http://bioinformatics.oxfordjournals.org/cgi/content/abstract/22/16/2047>.
- [40] Elena Rivas and Sean R. Eddy. Probabilistic phylogenetic inference with insertions and deletions. *PLoS Comput Biol*, 4(9):e1000172, 09 2008. doi: 10.1371/journal.pcbi.1000172. URL <http://dx.doi.org/10.1371%2Fjournal.pcbi.1000172>.
- [41] Kevin Liu, Sindhu Raghavan, Serita Nelesen, C. Randal Linder, and Tandy Warnow. Rapid and Accurate Large-Scale Coestimation of Sequence Alignments and Phylogenetic Trees. *Science*, 324(5934):1561–1564, 2009. doi: 10.1126/science.1171243. URL <http://www.sciencemag.org/cgi/content/abstract/324/5934/1561>.
- [42] Radu Mihaescu, Cameron Hill, and Satish Rao. Fast phylogeny reconstruction through learning of ancestral sequences. *CoRR*, abs/0812.1587, 2008.
- [43] Lusheng Wang and Tao Jiang. On the complexity of multiple sequence alignment. *Journal of Computa-*

- tional Biology*, 1(4):337–348, 1994.
- [44] Isaac Elias. Settling the intractability of multiple alignment. *Journal of Computational Biology*, 13(7):1323–1339, 2006. doi: 10.1089/cmb.2006.13.1323. URL <http://www.liebertonline.com/doi/abs/10.1089/cmb.2006.13.1323>. PMID: 17037961.
 - [45] D. G. Higgins and P. M. Sharp. Clustal: a package for performing multiple sequence alignment on a microcomputer. *Gene*, 73(1):237–244, 1988.
 - [46] C. Notredame, D.G. Higgins, and J. Heringa. T-coffee: A novel method for fast and accurate multiple sequence alignment. 2000.
 - [47] Kazutaka Katoh, Kazuharu Misawa, Kei-ichi Kuma, and Takashi Miyata. MAFFT: a novel method for rapid multiple sequence alignment based on fast Fourier transform. *Nucl. Acids Res.*, 30(14):3059–3066, 2002. doi: 10.1093/nar/gkf436. URL <http://nar.oxfordjournals.org/cgi/content/abstract/30/14/3059>.
 - [48] Robert C. Edgar. MUSCLE: multiple sequence alignment with high accuracy and high throughput. *Nucl. Acids Res.*, 32(5):1792–1797, 2004. doi: 10.1093/nar/gkh340. URL <http://nar.oxfordjournals.org/cgi/content/abstract/32/5/1792>.
 - [49] David Sankoff. Minimal mutation trees of sequences. *SIAM Journal on Applied Mathematics*, 28(1):35–42, 1975. doi: 10.1137/0128004. URL <http://link.aip.org/link/?SMM/28/35/1>.
 - [50] Lusheng Wang, Tao Jiang, and E. L. Lawler. Approximation algorithms for tree alignment with a given phylogeny. *Algorithmica*, 16(3):302–315, 1996.
 - [51] G. Navarro. A guided tour to approximate string matching. *ACM Computing Surveys (CSUR)*, 33(1):31–88, 2001.
 - [52] G. Navarro, R. Baeza-Yates, E. Sutinen, and J. Tarhio. Indexing methods for approximate string matching. *Bulletin of the Technical Committee on*, page 19.
 - [53] A. Andoni and R. Krauthgamer. The smoothed complexity of edit distance. *Lecture Notes in Computer Science*, 5125:357–369, 2008.
 - [54] Y. Higuchi. Remarks on the limiting Gibbs states on a $(d + 1)$ -tree. *Publ. Res. Inst. Math. Sci.*, 13(2):335–348, 1977.